

WHAT IS CLAIMED IS:

1. A system for providing one or more virtual volumes, comprising:
  - a host system;
  - a set of storage devices, each of which includes physical block addresses that stores data; and
  - a network switch system connecting the host system and the set of storage devices, and configured to define and manage a virtual volume associated with data distributed across the physical block addresses, the network switch system including:
    - a first virtualization layer that maintains first tier objects including information reflecting a relationship between the physical block addresses and one or more logical partitions of virtual volume data, and
    - a second virtualization layer that maintains second tier objects including information reflecting a logical configuration of the virtual volume,wherein the network switch system manages the virtual volume for the host system using the first and second tier objects.
2. The system of claim 1, wherein the host system sends to the network switch system a request to access the virtual volume and the network switch system leverages the first virtualization layer to collect information associated with the one or more logical partitions associated with the request.
3. The system of claim 1, wherein the second virtualization layer accesses the first virtualization layer using the second tier mappings based on a request received

from the host system to access information associated with the one or more logical partitions.

4. The system of claim 1, wherein the first virtualization layer includes first tier storage processors that are each selectively connected to one or more of the storage devices.

5. The system of claim 1, wherein the first tier mappings include state-based information associated with the logical partitions of the virtual volume.

6. The system of claim 1, wherein the logical configuration of the virtual volume is at least one of a mirrored virtual volume configuration, a concatenation configuration, a striped virtual volume configuration, and a striped over mirrored virtual volume configuration.

7. The system of claim 1, wherein the first virtualization layer includes first tier storage processors, each of which includes at least one of (i) a first tier mapping for the virtual volume and (ii) a second tier mapping.

8. The system of claim 7, wherein any first tier storage processor having a second tier mapping and a first tier mappings have communication access to the host system and at least one storage device.

9. The system of claim 1, wherein the second virtualization layer includes second tier storage processors, each of which includes one or more of the second tier mappings and each of which are connected to the host system.

10. The system of claim 1, wherein the network switch system includes a set of storage processors separated into first tier storage processors associated with the first virtualization layer and second tier storage processors associated with the second virtualization layer, and wherein the network switch system includes a switching fabric interconnecting the first tier storage processors and the second tier storage processors.

11. The system of claim 10, further including a second host system and wherein the each second tier storage processor is connected to at least one of the host system and the second host system, and wherein the set of storage devices are selectively connected to the first tier storage processors.

12. The system of claim 10, wherein each storage processor includes a virtualization state manager (VSM) that is configured to manage a local version of the virtual volume.

13. The system of claim 12, wherein each storage processor VSM is configured to manage any of the mappings maintained by the respective storage processor.

14. The system of claim 12, wherein each storage processor includes a Master VSM (MVSM) that is either in an active or inactive state.

15. The system of claim 14, wherein when activated, the MVSM is configured to interact with the VSMs of the storage processor to build a current system image of the virtual volume.

16. The system of claim 15, wherein the current system image of the virtual volume includes information reflecting which storage processors are connected to selective ones of the storage devices and which storage processors are connected to the host system.

17. The system of claim 14, wherein the network switch system designates a single storage processor as a Master Virtualization Storage Processor (MVSP) by activating the MVSM in the designated MVSP.

18. The system of claim 17, wherein the MVSP is configured to interface with a Virtualization State Manager DataBase (VSMDB) stored in the set of storage devices to build a system image of the virtual volume.

19. The system of claim 18, wherein the VSMDB is distributed across the storage devices such that each storage device includes a local VSMDB including

virtualization objects associated with any storage processors connected to the respective storage device.

20. The system of claim 19, wherein the MVSP collects the virtualization objects from each local VSMDB to build the system image.

21. The system of claim 20, wherein the MVSP sends the system image to a network switch system managing component to build a global system image reflecting a logical view of the virtual volume.

22. The system of claim 21, wherein the global system image includes first tier objects having object definitions associated with the logical partitions and second tier objects having object definitions associated with the logical configuration of the virtual volume and references to selective ones of the first tier objects.

23. The system of claim 22, wherein the network switch system managing component includes a Virtualization Coherency Manager (VCM) that maps the first tier objects to the first tier storage processors and maps the second tier objects to the second tier storage processors.

24. The system of claim 23, wherein the VCM distributes the first tier objects to the first tier storage processors before distributing the second tier objects.

25. The system of claim 24, wherein the VCM directs the first tier storage processors to expose their first tier objects to the second tier storage processors such that the second tier storage processors discover selective ones of the first tier objects.

26. The system of claim 25, wherein the second tier storage processors determine logical relationships between selective ones of the second tier objects and the first tier objects.

27. The system of claim 10, wherein the network switch system includes a Virtualization Coherency Manager (VCM) that manages the virtual volume using the first and second tier mappings.

28. The system of claim 27, wherein the VCM selects one of the storage processors as a Master Virtualization Storage Processor (MVSP) that is configured to build a current system image.

29. The system of claim 28, wherein the MVSP creates first and second tier object definitions corresponding to the virtual volume based on information collected from selected ones of the storage devices.

30. The system of claim 29, wherein the VCM assigns the first tier object definitions to selective ones of the first tier storage processors and the second tier object definitions to selective ones of the second tier storage processors.

31. A multi-tier switching system for interfacing at least one host system and a storage volume including a set of storage devices, the system comprising:

a set of storage devices storing data for a virtual volume associated with the at least one host system;

first tier storage processors, each of which is connected to at least one of the storage devices and includes first tier objects defining logical partitions associated with the virtual volume;

second tier storage processors, each of which is connected to the at least one host system and maintains second tier objects defining a logical configuration of the virtual volume ; and

a switching fabric interconnecting the first tier and second tier of storage processors,

wherein each second tier storage processor is configured to receive a request to access first data included in the virtual volume and to use a second tier object to access a first tier object associated with the first data using references defined by the second tier object, and wherein each first tier storage processor is configured to use a first tier object to identify a logical partition corresponding to the first data to extract the first data.

32. A method for providing one or more virtual volumes in a virtualization storage environment including a host system, a set of storage devices, each of which includes physical block addresses that stores data, and a network switch system connecting the host system and the set of storage devices, and configured to define and manage a virtual volume associated with data distributed across the physical block addresses, the method comprising:

providing a first virtualization layer that maintains first tier objects including information reflecting a relationship between the physical block addresses and one or more logical partitions of virtual volume data;

providing a second virtualization layer that maintains second tier objects including information reflecting a logical configuration of the virtual volume; and

managing, by the network switch system, the virtual volume for the host system using the first and second tier objects.

33. The method of claim 32, including:

sending, from the host system to the network switch system, a request to access the virtual volume; and

leveraging, by the network switch system, the first virtualization layer to collect information associated with the one or more logical partitions associated with the request.

34. The method of claim 32, further comprising:



accessing, by the second virtualization layer, the first virtualization layer using the second tier mappings based on a request received from the host system to access information associated with the one or more logical partitions.

35. The method of claim 32, wherein the first virtualization layer includes first tier storage processors that are each selectively connected to one or more of the storage devices.

36. The method of claim 32, wherein the first tier mappings include state-based information associated with the logical partitions of the virtual volume.

37. The method of claim 32, wherein the logical configuration of the virtual volume is at least one of a mirrored virtual volume configuration, a concatenation configuration, a striped virtual volume configuration, and a striped over mirrored virtual volume configuration.

38. The method of claim 32, wherein the first virtualization layer includes first tier storage processors, each of which includes at least one of (i) a first tier mapping for the virtual volume and (ii) a second tier mapping.

39. The method of claim 32, wherein each storage processor includes a virtualization state manager (VSM), and the method further includes:

managing, by each VSM, a local version of the virtual volume.

40. The method of claim 39, further including:

managing, by each storage processor VSM, any of the mappings maintained by the respective storage processor.

41. The method of claim 40, wherein each storage processor includes a

Master VSM (MVSM) that is either in an active or inactive state.

42. The method of claim 41, further including:

activating a selected storage processor's MVSM to allow the activated MVSM to interact with the VSMs of the storage processor to build a current system image of the virtual volume.

43. The method of claim 42, wherein the current system image of the virtual

volume includes information reflecting which storage processors are connected to selective ones of the storage devices and which storage processors are connected to the host system.

44. The method of claim 41, further including:

designating a single storage processor as a Master Virtualization Storage Processor (MVSP) by activating the MVSM in the designated MVSP.

45. The method of claim 44, further including:

interfacing, by the MVSP, with a Virtualization State Manager DataBase (VSMDB) stored in the set of storage devices to build a system image of the virtual volume.

46. The method of claim 32, wherein managing the virtual volume includes:  
processing requests from the host system to access or modify the virtual volume  
and processing requests from the host system to create a new virtual volume.